

ON THE EXPRESSIVE POWER OF OVERLAPPING OPERATIONS OF DEEP NETWORKS

Or Sharir & Amnon Shashua

The Hebrew University of Jerusalem

{or.sharir, shashua}@cs.huji.ac.il

ABSTRACT

Expressive Efficiency with respect to a network architectural attribute P refers to the property where an architecture without P must grow exponentially large in order to approximate the expressivity of a network with attribute P . For example, it is known that depth is an architectural attribute that generates exponential efficiency in the sense that a shallow network must grow exponentially large in order to approximate the functions represented by a deep network of polynomial size. In this paper we extend the study of expressive efficiency to the attribute of network connectivity and in particular to the effect of "overlaps" in the convolutional process, i.e., when the stride of the convolution is smaller than its kernel size (receptive field).

Our analysis shows that having overlapping local receptive fields, and more broadly denser connectivity, results in an exponential increase in the expressive capacity of neural networks. Moreover, while denser connectivity can increase the expressive capacity, we show that the most common types of modern architectures already exhibit exponential increase in expressivity, without relying on fully-connected layers.

1 INTRODUCTION

One of the most fundamental attributes of deep networks, and the reason for driving its empirical success, is the "Depth Efficiency" result which states that deeper models are exponentially more expressive than shallower models of similar size. Formal studies of Depth Efficiency include the early work on boolean or thresholded circuits (Sipser, 1983; Yao, 1989; Håstad and Goldmann, 1991; Hajnal et al., 1993), and the more recent studies covering the types of networks used in practice (Pascanu et al., 2013; Montúfar et al., 2014; Eldan and Shamir, 2016; Cohen et al., 2016a; Cohen and Shashua, 2016; Telgarsky, 2016; Safran and Shamir, 2016; Raghu et al., 2016; Poole et al., 2016).

What makes the Depth Efficiency attribute so desirable, is that it brings exponential increase in expressive power through merely a polynomial change in the model, i.e. the addition of more layers. Nevertheless, depth is merely one among many architectural attributes that define modern networks. The deep networks used in practice consist of architectural features defined by various schemes of connectivity, convolution kernel defined by size and stride, pooling geometry and activation functions. Whether or not those relate to expressive efficiency, as depth has proven to be, remains an open question.

In order to study the effect of network design on expressive efficiency we should first define "efficiency" in broader terms. Given two network architectures \mathcal{H}_1 and \mathcal{H}_2 defined by architectural property P , we say that architecture \mathcal{H}_1 is expressively P -efficient with respect to architecture \mathcal{H}_2 , if the following two conditions hold: (i) any function \mathbf{h} realized by \mathcal{H}_2 of size $|\mathbf{h}|$ can be realized (or approximated) by some function $\mathbf{h}^* \in \mathcal{H}_1$ with size $|\mathbf{h}^*| \in \mathcal{O}(|\mathbf{h}|)$; (ii) there exist a function \mathbf{h} realized by \mathcal{H}_1 that cannot be realized (or approximated) by $\mathbf{h}^* \in \mathcal{H}_2$ unless $|\mathbf{h}^*| \in \Omega(f(|\mathbf{h}|))$ for some super-linear function f . The exact definition of $|\mathbf{h}|$ depends on the measurement we care about, e.g. the number of parameters of \mathbf{h} , or the number of "neurons" in \mathbf{h} . The nature of the function f in condition (ii) determines the type of efficiency taking place – if f is exponential then architecture \mathcal{H}_1 is said to be exponentially P -efficient with respect to architecture \mathcal{H}_2 , and if f is

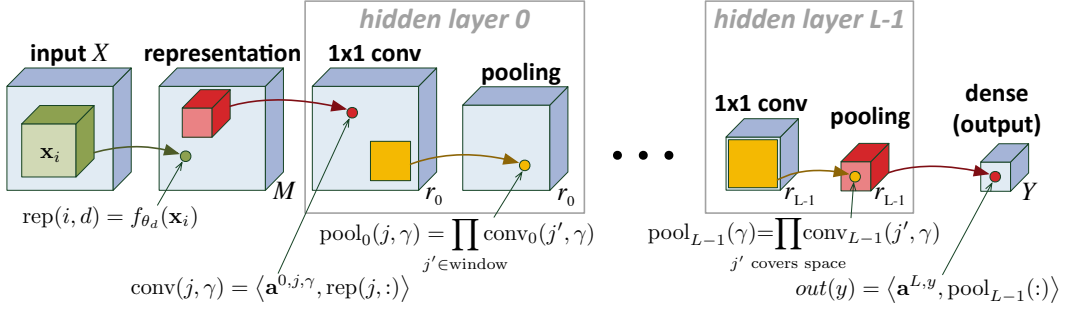


Figure 1: The original Convolutional Arithmetic Circuits as presented by Cohen et al. (2016a).

polynomial so is the expressive efficiency. Additionally, we say \mathcal{H}_1 is *completely more P-efficient* than \mathcal{H}_2 , if condition (ii) holds not just for some specific functions in \mathcal{H}_1 , but for all functions other than a negligible set.

In this paper we study the Efficiency associated with the architectural attribute of convolutions, namely the size of convolutional kernels (receptive field) and more importantly their stride. We say that a network architecture is of the *non-overlapping* type when the size of the local receptive field in each layer is equal to the stride. In that case, the sets of pixels participating in the computation of each two neurons in the same layer are completely separated. When the stride is smaller than the receptive field we say that the network architecture is of the *overlapping* type. In the latter case, the overlapping degree is determined by the *total* receptive field and stride projected back to the input layer — the implication being that for the overlapping architecture the total receptive field and stride can grow much faster than with the non-overlapping case.

We will derive a lower bound on the complete expressive capacity of the Overlapping architecture compared to the non-Overlapping. The lower bound depends on on the first layer for which its total receptive field exceeds a quarter of the input layer size. Since in the Overlapping architecture the total receptive field can grow fast, even with very small kernel size, the Efficiency becomes exponential for the majority of settings of kernel size and stride adopted in modern networks. Specifically, for a network comprising an alternating sequence of convolutional and pooling layers, with kernel size of 5×5 with unit stride followed by 2×2 pooling, yields a lower bound on Efficiency of M^{20} where M is the number of channels of the first hidden layer (typically, $M \approx 100$ in modern networks). In other words, for a non-Overlapping architecture to approximate the functions (up to a negligible set) realized by the Overlapping architecture above, one would require a network size whose lower bound behaves as 10^{40} .

2 PRELIMINARIES

We base our analysis on the convolutional arithmetic circuit (ConvAC) architecture introduced by Cohen et al. (2016a), which is illustrated by fig. 1, and can be simply thought of as a regular ConvNet, but with linear activations and product pooling layers, instead of the more common non-linear activations (e.g. ReLU) and average/max pooling. More specifically, each point in the input space of the network, denoted by $X = (\mathbf{x}_1, \dots, \mathbf{x}_N)$, is represented as an N -length sequence of s -dimensional vectors $\mathbf{x}_1, \dots, \mathbf{x}_N \in \mathbb{R}^s$. X is typically thought of as an image, where each \mathbf{x}_i corresponds to a local patche from that image. The first layer of the network is referred to as the representation layer, consisting of applying M representation functions $f_{\theta_1}, \dots, f_{\theta_M} : \mathbb{R}^s \rightarrow \mathbb{R}$ on each local patch \mathbf{x}_i , giving rise to M feature maps. Under the common setting, where the representation functions are selected to be $f_{\theta_d}(\mathbf{x}) = \sigma(\mathbf{w}_d^T \mathbf{x} + b_d)$ for some point-wise activation $\sigma(\cdot)$ and parameterized by $\theta_d = (\mathbf{w}_d, b_d) \in \mathbb{R}^s \times \mathbb{R}$, the representation layer reduces to the standard convolutional layer. Other possibilities, e.g. gaussian functions with diagonal covariances, have also been considered in Cohen et al. (2016a). Following the representation layer, are hidden layers indexed by $l = 0, \dots, L-1$, each begins with a 1×1 conv operator, which is just an $r_{l-1} \times 1 \times 1$ convolutional layer with r_{l-1} input channels and r_l output channels, with the sole exception that parameters of each kernel could be spatially unshared (known as locally-connected layer (Taigman et al., 2014)). Following each conv layer is a spatial pooling, that takes products of non-overlapping

two-dimensional windows covering the output of the previous layer, where for $l = L - 1$ the pooling window is the size of the entire spatial dimension (i.e. global pooling), reducing its output's shape to a $r_{L-1} \times 1 \times 1$, i.e. an r_{L-1} -dimensional vector. The final L layer maps this vector with a dense linear layer into the Y network outputs, denoted by $\mathbf{h}_y(\mathbf{x}_1, \dots, \mathbf{x}_N)$, representing score functions classifying each X to one of the classes through: $y^* = \operatorname{argmax}_y \mathbf{h}_y(\mathbf{x}_1, \dots, \mathbf{x}_N)$. As shown in Cohen et al. (2016a), these functions have the following form:

$$\mathbf{h}_y(\mathbf{x}_1, \dots, \mathbf{x}_N) = \sum_{d_1, \dots, d_N=1}^M \mathcal{A}_{d_1, \dots, d_N}^y \prod_{i=1}^N f_{\theta_{d_i}}(\mathbf{x}_i) \quad (1)$$

where \mathcal{A}^y , called the *coefficients tensor*, is a tensor of order N and dimension M in each mode, which for the sake of discussion can simply be seen as a multi-dimensional array, specified by N indices d_1, \dots, d_N each ranging in $\{1, \dots, M\}$, with entries given by polynomials in the network's *conv* weights. A byproduct of eq. 1 is that for a fixed set of M representation functions, all functions represented by ConvACs lay in the same subspace of functions.

The ConvAC framework has many desirable attributes. First, it shares many of the same traits as modern ConvNets, i.e. locality, sharing and pooling. Second, it can be shown to form a universal hypotheses space. Third, its underlying operations lend themselves to mathematical analysis based on measure theory and tensor analysis. Forth, through the concept of generalized tensor decompositions, ConvACs can be transformed to standard ConvNets with ReLU activations and average/max pooling, which laid the foundation for extending its proof methodologies to general ConvNets. Finally, from an empirical perspective, they tend to work well in many practical settings, e.g. for optimal classification with missing data (Sharir et al., 2016), and for compressed networks (Cohen et al., 2016b). The only limitation of the above framework is its underlying assumption that all convolutional layers are made strictly of 1×1 convolutions and non-overlapping pooling windows. In the next section we will introduce an extension of this framework to lift this limitation, allowing us to study the effects of overlaps on the expressivity of a network.

3 EXPRESSIVE CAPACITY OF OVERLAPPING OPERATIONS

In this section we analyze the expressive capacity of overlapping architectures against non-overlapping architectures. We begin by extending the ConvAC framework to support overlapping operations as described in sec. 3.1. We then proceed, in sec. 3.2, to state the preliminaries required to formalize our claims. With the preliminaries in place we present our main results in sec. 3.3, deriving both a generic lower-bound on the size of non-overlapping architectures, as well as examining special cases and their respective lower-bounds. Finally, we compare our results to the recent work by Cohen and Shashua (2017) on "Pooling Geometry" as described in sec. 3.4.

3.1 GENERALIZED CONVOLUTIONAL LAYER

As mentioned in the previous section, while ConvACs offer a rich theoretical foundation to study convolutional networks, they do not support overlapping operations — a limitation we aim to lift in this section. We accomplish this by generalizing the hidden layers of the ConvAC model into a unifying convolutional layer that supports overlaps, such that under non-overlapping settings it reduces to the previously defined hidden layer, comprising a pair of 1×1 *conv* layer and product pooling layer.

We begin by defining the Generalized Convolutional Operation (GCO) for $\mathbf{x} \in \mathbb{R}^{D \times R^2}$, parameterized by weights $\mathbf{w} \in \mathbb{R}^{D \times R^2}$ and biases $\mathbf{b} \in \mathbb{R}^{R^2}$, as follows:

$$\text{GCO}(\mathbf{x}; \mathbf{w}, \mathbf{b}) = \prod_{k=1}^{R^2} \left(b_k + \sum_{d=1}^D w_{dk} x_{dk} \right) \quad (2)$$

Given the above definition, we define the Generalized Convolutional (GC) layer as the convolutional application of the GCO on patches from the input to the layer, as illustrated in fig. 2. More specifically, assuming the input is of width and height equal to W and has D channels, a GC layer with local receptive field $R \times R$, strides $S \times S$, and C output channels, generates an output array




Figure 2: Illustration of the Generalized Convolutional (GC) layer with C output channels, local receptive field $R \times R$, and strides $S \times S$: **(a)** The GC layer operate by extracting from the input N^2 patches of size $R \times R$, where $N = \lceil W/S \rceil$, followed by applying on each patch $\mathbf{x}^{(j)}$ the GCO operation, parameterized by $(\mathbf{w}^{(c)}, \mathbf{b}^{(c)})$ for each output channel $c \in [C]$. **(b)** The application of GCO on patches can be thought of as the application of a ConvAC sub-network, comprising 1×1 conv layer followed by global product pooling, similar in principle to the “Network in Network” architecture (Lin et al., 2014).




Figure 3: The same ConvAC architecture from fig. 1, where each hidden layer is replaced with the newly defined GC layer, while constraining its local receptive field to be equal to its stride.

of width and height equal to $N = \lceil W/S \rceil$ and C channels, where each entry of the array corresponds to the result of applying the GCO on a patch $\mathbf{x}^{(j)} \in \mathbb{R}^{D \times R^2}$ and a pair of weights and biases $(\mathbf{w}^{(c)} \in \mathbb{R}^{D \times R^2}, \mathbf{b}^{(c)} \in \mathbb{R}^{R^2})$ for $j \in [N^2]$ and $c \in [C]$. For simplicity, we assume patches stepping out of bounds encounter zero padding. Notice that the GC layer can be thought of as convolving a ConvAC sub-network over its input, comprising a 1×1 conv layer followed by global product pooling, which bear resemblance to the “Network in Network” architecture (Lin et al., 2014). Additionally, just as the conv layer of the original ConvAC that was described in sec. 2, the parameters of the GC layer could optionally be unshared, i.e. for every patch $\mathbf{x}^{(j)}$ there would be a separate set of weights $\{(\mathbf{w}^{(c,j)}, \mathbf{b}^{(c,j)})\}_{c \in [C]}$, however, for the majority of the article we assume this is not the case.

Under the non-overlapping settings of $R = S$, i.e. when the local receptive field equals to the stride, the GC layer reduces to a 1×1 conv layer followed by a product pooling layer with $R \times R$ pooling windows. If the parameters of the GC layer are unshared, then so are the induced parameters of the corresponding 1×1 conv layer. In the more common shared parameters case, the induced parameters of the corresponding 1×1 conv layer are partially shared, following a repeating grid pattern of weights. With the above correspondence between non-overlapping GC layer and the previously defined hidden layer of the original ConvAC framework, we can redefine the ConvAC architecture depicted in fig. 1 using the newly defined GC layers, as illustrated in fig. 3.

While we have just established that the non-overlapping setting of the GC layer is equivalent to vanilla ConvACs, one might wonder if using overlapping layers instead could diminish what these overlapping networks can represent. We show that not only is it not the case, but prove the more general claim that a network can always realize a similar network with smaller local receptive fields, including the non-overlapping case where the local receptive fields equal exactly their stride.

Proposition 1. Let Φ and Ψ be two networks comprising a sequence of GC layers. If the architecture of Ψ can be derived from Φ through the removal of layers with 1×1 stride, or by decreasing the

local receptive field of some of the layers, then for any setting of parameters of Ψ , there exists a matching setting of parameters for Φ , such that the function realized by Ψ is exactly equivalent to Φ . Specifically, Φ can realize any non-overlapping network with the same order of strides (excluding 1×1 strides).

Proof. (Sketch) This follows from two simple claims: (i) A GC layer can produce an output equivalent to that of a GC layer with a smaller local receptive field, by “zeroing” its weights beyond the smaller local receptive field; and (ii) GC layers can be set such that their output is equal to their input, i.e. realize the identity function. With these claims, the local receptive fields of Φ can be effectively shrunk to match the local receptive fields of Ψ , and any additional layers of Φ with stride 1×1 could be set such that they are realizing the identity mapping, effectively “removing” them from Φ . See app. A.2 for the complete proof. \square

Proposition 1 essentially means that networks with overlapping operations do not decrease the expressivity of the network compared to networks of similar structure, i.e. same order of strides that are greater than 1×1 , but without overlaps. As we recall, this satisfies the first condition of the efficiency property introduced in sec. 1, regardless if we measure the size of a network as the number of parameters, or the number of “neurons”¹. In the following section we will cover the preliminaries required to show that overlapping networks actually lead to an increase in expressive capacity, which under some setting results in exponential gain, proving that the second condition of the efficiency property holds for the overlaps attribute.

3.2 SETTING UP THE NOTATIONS

In this section we present the preliminaries required to state our main results. To begin with, we wish to compare non-overlapping ConvACs to overlapping ConvACs, for a fixed set of M representation functions. While all functions realizable by non-overlapping ConvACs with shared representation functions lay in the same function subspace, this is not the case for overlapping ConvACs, which can realize additional functions outside the sub-space induced by non-overlapping ConvACs. We cannot therefore compare both architectures directly, and need to compare them through an auxiliary objective. Following the work of Cohen and Shashua (2016), we instead compare architectures through the concept discretization known as “grid tensors”. The underlying idea is to fix a set of points on an exponentially large grid in the input space, and examine the possible values each architecture can attain over these set of inputs. Specifically, fixing a set of *template* vectors $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(M)} \in \mathbb{R}^s$, the points on the grid are the set $\{(\mathbf{x}^{(d_1)}, \dots, \mathbf{x}^{(d_N)})\}_{d_1, \dots, d_N=1}^M$, and for a given function $\mathbf{h}(\mathbf{x}_1, \dots, \mathbf{x}_N)$ its set of values arranged in the form of a tensor (multi-dimensional array) are $\mathcal{A}(\mathbf{h})_{d_1, \dots, d_N} = \mathbf{h}(\mathbf{x}^{(d_1)}, \dots, \mathbf{x}^{(d_N)})$, called the grid tensor induced by \mathbf{h} . As proven in Cohen and Shashua (2016), assuming the fixed representation functions are linearly independent, then there exists templates $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(M)}$, such that any non-overlapping ConvAC can represent all possible grid tensors over these templates, given sufficient number of channels at each layer. Thus, once we fix a set of linearly independent representation functions, we can compare different ConvACs, whether overlapping or not, on the minimal size required for them to induce the same grid tensor, while knowing such a finite number always exists.

With the foundations to compare between overlapping and non-overlapping networks in place, we move on to defining the characterizing properties controlling the expressivity of overlapping networks. Let Φ be a ConvAC with M linearly independent representation functions, comprising L GC layers, where for all $l \in [L]$ the l ’th layer has a local receptive field of $R^{(l)} \times R^{(l)}$, a stride of $S^{(l)} \times S^{(l)}$, and $D^{(l)}$ output channels. We define the *total stride* of the l ’th layer, denoted by $T_S^{(l)} := T_S(S^{(1)}, \dots, S^{(l)})$, and the *total receptive field* of the l ’th layer, denoted by

¹We take here the broader definition of a “neuron”, as any one of the scalar values comprising the output array of an arbitrary layer in a network. In the case the output array is of width and height equal to N and C channels, then the number of such “neurons” for that layer is $N^2 C$.

$T_R^{(l)} := T_R(R^{(1)}, S^{(1)}, \dots, R^{(l)}, S^{(l)})$, as follows:

$$T_S^{(l)}(S^{(1)}, \dots, S^{(l)}) = \begin{cases} \prod_{i=1}^l S^{(i)} & l \geq 1 \\ 1 & l = 0 \end{cases} \quad (3)$$

$$T_R^{(l)}(R^{(1)}, S^{(1)}, \dots, R^{(l)}, S^{(l)}) = R^{(l)} \cdot T_S^{(l-1)} + \sum_{k=1}^{l-1} (R^{(k)} - S^{(k)}) \cdot T_S^{(k-1)} \quad (4)$$

The total receptive field and stride are best thought of as the receptive field and stride with respect to the input layer. Considering proposition 1, ConvACs can also realize effectively smaller local receptive fields, which raise the question of what is the minimal total receptive field realizable by a network that is still greater than some value α . We define this value as the α -minimal total receptive field, and denote it by $T_R^{(l, \alpha)}$, which equals to:

$$T_R^{(l, \alpha)} = \underset{\substack{\forall i \in [l], S^{(i)} \leq t_i \leq R^{(i)} \\ T_R^{(l)}(t_1, S^{(1)}, \dots, t_l, S^{(l)}) > \alpha}}{\operatorname{argmin}} T_R^{(l)}(t_1, S^{(1)}, \dots, t_l, S^{(l)}) \quad (5)$$

Notice that for non-overlapping networks the total receptive field always equals the total stride, and that only at the end of the network, after the spatial dimension collapses to 1×1 , does the total receptive field grow to encompass the entire size of the representation layer. For overlapping networks this is not the case, and the total receptive field could grow much faster. Intuitively, this means that values in regions of the input layer that are far apart would be combined by non-overlapping networks only near the last layers of network, and thus non-overlapping networks are effectively shallow in comparison to overlapping networks. Based on this intuition, in the next section we analyze networks with respect to the point at which their total receptive field is large enough, i.e., above a certain threshold.

3.3 MAIN RESULTS

With all the preliminaries in place, we are ready to present our main result presented below:

Theorem 1. *Let Φ be a ConvAC as described above, where the output of its representation layer is of width and height equal to W and has M channels. Following the representation layer are L GC layers, where the l 'th layer has a local receptive field of $R^{(l)} \times R^{(l)}$, a stride of $S^{(l)} \times S^{(l)}$, and $D^{(l)}$ output channels. Let \tilde{L} be a layer such that its total receptive field is greater than a quarter of the output of the representation layer, i.e. that $T_R^{(\tilde{L})}(R^{(1)}, S^{(1)}, \dots, R^{(\tilde{L})}, S^{(\tilde{L})}) > \frac{W}{2}$, and define $\bar{D} = \min\{M, D^{(\tilde{L})}, \frac{1}{2} \min_{1 \leq l < \tilde{L}} D^{(l)}\}$. Suppose we randomize the parameters of Φ by some continuous distribution. Then, with probability 1, we obtain a function that cannot be realized (or approximated) by a non-overlapping ConvAC if the number of channels in its next to last layer is less than:*

$$\bar{D} \left\lceil \frac{W - T_R^{(\tilde{L}, \lceil W/2 \rceil)}}{T_S^{(\tilde{L})}} + 1 \right\rceil \cdot \left\lceil \frac{W}{T_S^{(\tilde{L})}} \right\rceil \quad (6)$$

Proof. (Sketch) Given the structure of an overlapping network, we wish to show that with probability 1, we obtain a function that cannot be realized by a non-overlapping network, unless the number of channels of its next to last layer is above a certain threshold. As discussed in sec. 3.2, instead of finding a lower bound for realizing a function, we find a lower bound for inducing the same grid tensor.

Our proof is built on the concept of tensor matricization, i.e. reordering the entries of a tensor to the shape of a matrix. The exact ordering of the entries is specified through the partition of the indices of a tensor into two even disjointed sets, where one set represent the rows and the other the columns of the resulted matrix. Recall that each index of the grid tensor relates to a spatial location of the input, which bring forth two natural partitions: “left-right” and “top-bottom” even partitions.

Through methods of tensor analysis, we adapt corollaries from Cohen and Shashua (2017), and show that the matricization rank, according to either of these matricizations of the induced grid tensor, serves as a lower bound on the size of non-overlapping networks. Thus, it is enough to show




Figure 4: A network architectures beginning with large local receptive fields greater than $N/2$ and at least M output channels. According to theorem 1, if we randomize the weights of this network, then almost always we obtain a function that cannot be approximated by a non-overlapping network if the number of channels in its next to last layer is less than $M^{\frac{W^2}{2}}$.

that the matricization rank of grid tensors induced by an overlapping network of a given structure are greater than the bound given by eq. 6.

In order for us to prove that, when we randomize the parameters of the network, the matricization rank holds the above inequality with probability 1, is equivalent to proving that the set of parameters for which this is not the case is of measure zero, with respect to the Lebesgue measure over the Euclidean space of the parameters. Recall that the entries of the grid tensor induced by an overlapping network are polynomial functions of its parameters, and according to a lemma from Sharir et al. (2016), if we find just a single example for which the matricization rank is greater than the desired lower bound, then this is true almost everywhere, i.e. with probability 1.

Given the last remark, the central part of our proof is simply the construction of such an example. First we find a set of parameters for the simpler case where the first GC layer is greater than a quarter of the input, satisfying the conditions of the theorem. The motivation behind the specific construction is the pairing of indices from each side of the partition, such that they are both in the same local receptive field, and designing the kernels such that the output of each local application of them defines a mostly diagonal matrix of rank \bar{D} , with respect to these two indices. The rest of the parameters are chosen such that the output of the entire network results in a product of the entries of these matrices. Under matricization, this results in a matrix who is equivalent² to a Kronecker product of mostly diagonal matrices. Thus, the matricization rank is equal to the product of the ranks of these matrices, which results in the exponential form of eq. 6.

Finally, we extend the above example to the general case, by realizing the operation of the first layer of the above example through multiple layers with small local receptive fields. See app. A.1 for the exact definitions and lemmas we base our proof on, and see app. A.3 for the complete proof.

□

While the complexity of the generic lower-bound above might seem incomprehensible at first, its generality gives us the tools to analyze practically any kind of feed-forward architecture. As an example, we can analyze the lower bound for the well known GoogLeNet architecture (Szegedy et al., 2015), for which the lower bound equals 32^{98} , making it clear that using a non-overlapping architecture for this case is infeasible. Next, we will see how theorem 1 can help us derive lower bounds for commonly used networks.

According to theorem 1, the lower bound depends on the first layer for which its total receptive field is greater than a quarter of the input. As mentioned in the previous section, for non-overlapping networks this layer is always the first layer \tilde{L} for which the spatial dimension collapses to 1×1 , which entails the total receptive field and total stride to be equal to the width of the representation layer, i.e. $T_R^{(\tilde{L})} = T_S^{(\tilde{L})} = W$, and substituting this values in eq. 6 results simply in \bar{D} – a trivial result that to realize a non-overlapping network the next to last layer must have at least half the channels of the target network.

²Two matrices are equivalent if they could be converted into one another through elementary row or column operations.




Figure 5: The common network architecture of alternating $B \times B$ “conv” and 2×2 “pooling” layers. If $B \leq W/5 + 1$ and $r_l \geq 2M$ for all $0 \leq l < L$, then the lower bound of theorem 1 for this network results in $M^{\frac{(2B-1)^2}{4}}$.

On the other extreme, we can examine the case where the first GC layer has a local receptive field R greater than a quarter of its input, i.e. $R > W/2$. Since the layers following the first GC layer do not affect the lower bound in this case, it applies to any arbitrary sequence of layers as illustrated in fig. 4. For simplicity we will also assume that the stride S is less than $W/2$, and that $\frac{W}{2}$ is evenly divided by S . In this case the $\frac{W}{2}$ -minimal receptive field equals to $\frac{W}{2} + 1$, and thus the lower bound results in $\bar{D}^{\frac{W^2}{2S}}$. Consider the case of $\bar{D} = M$ and $S = 1$, then a non-overlapping network satisfying this lower bound is of the order of magnitude in which it could already represent any possible grid tensor. This demonstrate our point from the introduction, that through a a polynomial change in the architecture, i.e. increasing the receptive field, we get an exponential increase in expressive capacity.

Though our previous example demonstrates that a polynomially sized overlapping network could lead to the maximal achievable lower-bound, in practice, applying such large convolutions is a very resource intensive task. The common best practice is to use multiple small local receptive fields of size $B \times B$, where the typical values are $B = 3$ or $B = 5$, separated by a 2×2 “pooling” layers, i.e. layers with both stride and local receptive field equal to 2×2 . For simplicity, we assume that $W = 2^L$ for some $L \in \mathbb{N}$. See fig. 5 for an illustration of such a network. Analyzing the above network with theorem 1 results in the following proposition.

Proposition 2. Consider a network comprising a sequence of GC blocks, each block begins with a layer whose local receptive field is $B \times B$ and its stride 1×1 , followed by a layer with local receptive field 2×2 and stride 2×2 , where the output channels of all layers are at least $2M$, and the spatial dimension of the representation layer is $W \times W$ for $W = 2^L$. Then, the lower bound describe by eq. 6 for the above network is greater than or equal to

$$\tau(B, W) := M^{\frac{(2B-1)^2}{2} \cdot (1 + \frac{2B-2}{W})^{-2}} = M^{\frac{W^2}{2} \cdot (1 + \frac{W-1}{2B-1})^{-2}}$$

whose limits are $\lim_{B \rightarrow \infty} \tau(B, W) = M^{\frac{W^2}{2}}$ and $\lim_{W \rightarrow \infty} \tau(B, W) = M^{\frac{(2B-1)^2}{2}}$. Assuming $B \leq \frac{W}{5} + 1$, then $\tau(B, W) \geq M^{\frac{(2B-1)^2}{4}}$.

Proof. (Sketch) We first find a closed-form expression for both the total stride and total receptive field for each of the $B \times B$ layers in the given network. We then show that for layers whose total receptive field is greater than $\frac{W}{2}$, its α -minimal total receptive field, for $\alpha = \frac{W}{2}$, is equal to $\frac{W}{2} + 1$. We then use the above to find the first layer who satisfies the conditions of theorem 1, and then use our closed-forms expressions to simplify the general lower bound for this specific network. See app. A.4 for the complete proof. \square

In particular, for the typical values of $M = 64$, $B = 5$, and $W \geq 20$, the lower bound is at least 64^{20} , which once again entails such networks could not be realized efficiently by non-overlapping networks. When B grows in size, this bound approaches the earlier result we have shown for large local receptive fields encompassing more than a quarter of the image. When W grows in size, the lower bound is dominated strictly by the local receptive fields. Also notice that based on proposition 2, we could also derive a respective lower bound for a network following VGG style architecture (Simonyan and Zisserman, 2014), where instead of a single convolutional layer before every “pooling” layer, we have K layers, each with a local receptive field of $C \times C$. Under this case, it is trivial to show that the bound from proposition 2 holds for $B = K \cdot (C - 1) + 1$. Typically values

of $C = 3$ and $K = 2$ are used, which result in an effective value of $B = 2 \cdot (3 - 1) + 1 = 5$, just as in the previous case.

3.4 COMPARISON TO POOLING GEOMETRY

From theorem 1 we learn that overlaps give rise to networks which almost always cannot be efficiently implemented by non-overlapping ConvAC with standard pooling geometry. However, as proven by Cohen and Shashua (2017), a ConvAC that uses a different pooling geometry – i.e. the input to the pooling layers are not strictly contiguous windows from the previous layer – also cannot be efficiently implemented by the standard ConvAC with standard pooling geometry. This raises the question of whether overlapping operations are simply equivalent to a ConvAC with a different pooling geometry and nothing more. We answer this question in two parts. First, a ConvAC with a different pooling geometry might be able to implement some function more efficiently than ConvAC with standard pooling geometry, however, the reverse is also true, that a ConvAC with standard pooling can implement some functions more efficiently than ConvAC with alternative pooling. In contrast, a ConvAC that uses overlaps is still capable to implement efficiently any function that a non-overlapping ConvAC with standard pooling can. Second, in theorem 2 below we show that it is possible under some architectures to efficiently realize functions that no non-overlapping ConvAC can realize efficiently, regardless of its pooling geometry.

Theorem 2. *Under the same settings as before, consider a network whose representation layer is followed by a GC layer with local receptive field $W \times W$, stride 1×1 , and $D \geq M$ output channels, whose parameters are “unshared” (see sec 3.1), followed by $(L - 1)$ arbitrary GC layers, whose final output is a scalar. Suppose we randomize the parameters of this network by some continuous distribution. Then, with probability 1, we obtain a function that cannot be realized (or approximated) by a non-overlapping ConvAC, regardless of its pooling geometry, if the number of channels in its next to last layer is less than $M^{\frac{W^2}{2}}$. The exact same result holds if the parameters of the first GC layers are “shared” and $D \geq M \cdot W^2$.*

Proof. (Sketch) We follow the same steps of our proof of theorem 1, utilizing the fact that for each pooling geometry there exists a matricization of the grid tensor, such that its matricization rank is a lower bound on the number of channels in its next to last layer. However, we do not construct just one specific overlapping network that attains a rank of $D \geq M \cdot W^2$, for all possible matricizations of the induced grid tensor. Instead, we construct a separate network for each possible matricization. This proves that with respect to the Lebesgue measure over the network’s parameters space, separately for each pooling geometry, the set of parameters for which the lower bound does not hold is of measure zero. Since a finite union of zero measured sets is also of measure zero, then the lower bound with respect to all possible pooling geometries holds almost everywhere, which concludes the proof sketch. See app. A.5 for the complete proof. \square

It is important to note that though the above theorem shows that pooling geometry on its own is less expressive than overlapping networks with standard pooling, it does not mean that pooling geometry is irrelevant. Specifically, we do not yet know the effect of combining both overlaps and alternative pooling geometries together. Additionally, many times sufficient expressivity is not the main obstacle for solving a specific task, and the inductive bias induced by a carefully chosen pooling geometry could help reduce overfitting.

4 DISCUSSION

The common belief amongst deep learning researchers has been that depth is one of the key factors in the success of deep neural networks, and that it should be prioritized over dense connectivity. Indeed, the contemporary best practices of neural network design advise that the majority of the layers in the network should be of the convolutional type, and only few if any fully-connected layers (Simonyan and Zisserman, 2014; Springenberg et al., 2014; He et al., 2016). Many speculate about the regularization properties of this practice, however, to date none have analyzed its exact effect on the expressive capacity. Cohen et al. (2016a) have already shown that limiting networks to just 1×1 convolutions and 2×2 pooling, under which the networks are non-overlapping, does not hinder the universality property of ConvNets, which given the aforementioned trend in network

design, raises the question of why overlapping architectures have been an integral part of all known successful deep learning models.

Our analysis shows that having overlapping local receptive fields, and more broadly denser connectivity, results in an exponential increase in the expressive capacity of neural networks, proving our Overlaps Efficiency conjecture we have introduced in sec. 1. Moreover, while denser connectivity can increase the expressive capacity, we show that the most common types of modern architectures already exhibit exponential increase in expressivity, without relying on fully-connected layers. This could partly explain the somewhat surprising observation, that networks comprised of mainly convolutional layers with small local receptive fields typically outperform networks employing larger local receptive fields and several fully-connected layers. Since the former networks are already in the exponential regime of expressivity, it is probable they are sufficiently expressive for most practical needs, and any further increase in connectivity would either barely affect the performance, or lead to overfitting. Several empirical works (Li and Perona, 2005; Coates et al., 2011) have demonstrated similar behavior, showing that the classification accuracy of networks can sharply decline as the degree of overlaps is decreased, while also showing that gains from using very large local receptive fields are insignificant compared to the increase in computational resources. Other works studying the receptive fields of neural networks have mainly focused on how to learn them from the data (Coates and Ng, 2011; Jia et al., 2012). While our analysis has no direct implications to those specific works, it does lay the ground work for potentially guiding architecture design, through quantifying the expressivity of any given architecture.

In theorem 1, we prove that the expressivity of the network is dependent on two factors, the total stride and total receptive field. Luo et al. (2016) study the *effective total receptive field* of different layers, where they measure the degree to which each input pixel is affecting the output of each activation. They show that under common random initialization of the weights, the effective total receptive field has a gaussian shape and is much smaller than the maximal total receptive field. They additionally demonstrate that during training the effective total receptive field grows in size, and suggests that weights should be initialized such that the initial effective receptive field is large. Their results strengthen our theory, by showing that trained networks tend to maximize their effective receptive field, taking full potential of their expressive capacity.

An interesting outcome, of moving from non-overlapping architectures to overlapping ones, is that the depth of a network is no longer capped at $\log_2(\text{input size})$, as has been the case in the models investigated in Cohen et al. (2016a). With our current theorems we cannot assess the contribution of depths beyond this limit, as we measure the expressivity of overlapping architecture only with respect to non-overlapping architecture. A natural extension of our current theory would be the analysis of upper bounds in addition to the lower bounds we have already discovered, which will allow us to compare between different overlapping networks. We suspect that networks whose depth goes beyond the aforementioned limit will play an important role in such analysis.

Finally, our current analysis is grounded on the framework of ConvACs, which were shown to produce very promising results in special applications, e.g. under missing data (Sharir et al., 2016) and for compressed networks (Cohen et al., 2016b). We have also carried out limited experiments with the specific models presented in this article, finding comparable performance to conventional ConvNets on both MNIST and CIFAR10 – the results of which will be published in a separate article. Though our theoretical results are based on ConvACs, other studies (Cohen and Shashua, 2016) have already shown that results on ConvACs could be transfer to the domain of conventional ConvNets using most of the same mathematical machinery presented in app A. The adaptation of our analysis according to the latter is left for future work.

ACKNOWLEDGMENTS

This work is supported by Intel grant ICRI-CI #9-2012-6133, by ISF Center grant 1790/12 and by the European Research Council (TheoryDL project).

REFERENCES

Adam Coates and Andrew Y Ng. Selecting Receptive Fields in Deep Networks. *Advances in Neural Information Processing Systems*, 2011.

- Adam Coates, Andrew Y Ng, and Honglak Lee. An Analysis of Single-Layer Networks in Unsupervised Feature Learning. *International Conference on Artificial Intelligence and Statistics*, pages 215–223, 2011.
- Nadav Cohen and Amnon Shashua. Convolutional Rectifier Networks as Generalized Tensor Decompositions. In *International Conference on Machine Learning*, May 2016.
- Nadav Cohen and Amnon Shashua. Inductive Bias of Deep Convolutional Networks through Pooling Geometry. In *International Conference on Learning Representations ICLR*, April 2017.
- Nadav Cohen, Or Sharir, and Amnon Shashua. On the Expressive Power of Deep Learning: A Tensor Analysis. In *Conference on Learning Theory COLT*, May 2016a.
- Nadav Cohen, Or Sharir, and Amnon Shashua. Deep SimNets. In *Computer Vision and Pattern Recognition CVPR*, May 2016b.
- Ronen Eldan and Ohad Shamir. The Power of Depth for Feedforward Neural Networks. In *Conference on Learning Theory COLT*, May 2016.
- Wolfgang Hackbusch. *Tensor Spaces and Numerical Tensor Calculus*, volume 42 of *Springer Series in Computational Mathematics*. Springer Science & Business Media, Berlin, Heidelberg, February 2012.
- András Hajnal, Wolfgang Maass, Pavel Pudlák, Mario Szegedy, and György Turán. Threshold Circuits of Bounded Depth. *J. Comput. Syst. Sci. ()*, 46(2):129–154, 1993.
- Johan Håstad and Mikael Goldmann. On the Power of Small-Depth Threshold Circuits. *Computational Complexity ()*, 1(2):113–129, 1991.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. *Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- Yangqing Jia, Chang Huang, and Trevor Darrell. Beyond spatial pyramids - Receptive field learning for pooled image features. *Computer Vision and Pattern Recognition*, 2012.
- Fei-Fei Li and Pietro Perona. A Bayesian Hierarchical Model for Learning Natural Scene Categories. *Computer Vision and Pattern Recognition*, 2:524–531, 2005.
- Min Lin, Qiang Chen, and Shuicheng Yan. Network In Network. In *International Conference on Learning Representations ICLR*, 2014.
- Wenjie Luo, Yujia Li, Raquel Urtasun, and Richard S Zemel. Understanding the Effective Receptive Field in Deep Convolutional Neural Networks. *Advances in Neural Information Processing Systems*, pages 4898–4906, 2016.
- Guido F Montúfar, Razvan Pascanu, Kyunghyun Cho, and Yoshua Bengio. On the Number of Linear Regions of Deep Neural Networks. *CoRR abs/1202.2745*, stat.ML:2924–2932, 2014.
- Razvan Pascanu, Guido Montufar, and Yoshua Bengio. On the number of response regions of deep feed forward networks with piece-wise linear activations. *arXiv.org*, December 2013.
- Ben Poole, Subhaneil Lahiri, Maithreyi Raghu, Jascha Sohl-Dickstein, and Surya Ganguli. Exponential expressivity in deep neural networks through transient chaos. In *Advances in Neural Information Processing Systems NIPS, Deep Learning Workshop*, 2016.
- Maithra Raghu, Ben Poole, Jon Kleinberg, Surya Ganguli, and Jascha Sohl-Dickstein. On the expressive power of deep neural networks. *arXiv.org*, June 2016.
- Itay Safran and Ohad Shamir. Depth Separation in ReLU Networks for Approximating Smooth Non-Linear Functions. *CoRR abs/1202.2745*, cs.LG, 2016.
- Or Sharir, Ronen Tamari, Nadav Cohen, and Amnon Shashua. Tractable Generative Convolutional Arithmetic Circuits. *arXiv.org*, October 2016.
- Karen Simonyan and Andrew Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR abs/1202.2745*, cs.CV, 2014.
- Michael Sipser. *Borel sets and circuit complexity*. ACM, New York, New York, USA, December 1983.
- Jost Tobias Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin Riedmiller. Striving for Simplicity: The All Convolutional Net. *arXiv.org*, December 2014.

Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going Deeper with Convolutions. In *Computer Vision and Pattern Recognition CVPR*, 2015.

Yaniv Taigman, Ming Yang, Marc’Aurelio Ranzato, and Lior Wolf. DeepFace: Closing the Gap to Human-Level Performance in Face Verification. In *Computer Vision and Pattern Recognition CVPR*. IEEE Computer Society, June 2014.

Matus Telgarsky. Benefits of depth in neural networks. In *Conference on Learning Theory COLT*, May 2016.

Andrew Chi-Chih Yao. Circuits and Local Computation. *STOC*, pages 186–196, 1989.

A DEFERRED PROOFS

In this section we present our proofs for the theorems and claims stated in the body of the article.

A.1 PRELIMINARIES

In this section we lay out the preliminaries required to understand the proofs in the following sections. We begin with a limited introduction to tensor analysis, followed by quoting a few relevant known results relating tensors to ConvACs.

We begin with basic definitions and operations relating to tensors. Let $\mathcal{A} \in \mathbb{R}^{M_1 \otimes \dots \otimes M_N}$ be a tensor of order N and dimension M_i in each mode $i \in [N]$ (where $[N] = \{1, \dots, N\}$), i.e. $\mathcal{A}_{d_1, \dots, d_N} \in \mathbb{R}$ for all $i \in [N]$ and $d_i \in [M_i]$. For tensors $\mathcal{A}^{(1)}$ and $\mathcal{A}^{(2)}$ of orders $N^{(1)}$ and $N^{(2)}$, and dimensions $M_{i_1}^{(1)}$ and $M_{i_2}^{(2)}$ in each of the modes $i_1 \in [N^{(1)}]$ and $i_2 \in [N^{(2)}]$, respectively, we define their tensor product $\mathcal{A}^{(1)} \otimes \mathcal{A}^{(2)}$ as the order $N^{(1)} + N^{(2)}$ tensor, where

$$\left(\mathcal{A}^{(1)} \otimes \mathcal{A}^{(2)} \right)_{d_1, \dots, d_{N^{(1)}+N^{(2)}}} = \mathcal{A}_{d_1, \dots, d_{N^{(1)}}}^{(1)} \cdot \mathcal{A}_{d_{N^{(1)}+1}, \dots, d_{N^{(1)}+N^{(2)}}}^{(2)}$$

For a set of vectors $\mathbf{v}^{(1)} \in \mathbb{R}^{M_1}, \dots, \mathbf{v}^{(N)} \in \mathbb{R}^{M_N}$, the N ordered tensor $\mathcal{A} = \mathbf{v}^{(1)} \otimes \dots \otimes \mathbf{v}^{(N)}$ is called an *elementary tensor*, or *rank-1 tensor*. More generally, any tensor can be represented as a linear combination of rank-1 tensors, i.e. $\mathcal{A} = \sum_{z=1}^Z \mathbf{v}^{(Z,1)} \otimes \dots \otimes \mathbf{v}^{(Z,N)}$, known as *rank-1 decomposition*, or CP decomposition, where the minimal Z for which this equality holds is known as the *tensor rank* of \mathcal{A} . Given a set of matrices $F^{(1)} \in \mathbb{R}^{M'_1 \times M_1}, \dots, F^{(N)} \in \mathbb{R}^{M'_N \times M_N}$, we denote by $\mathbf{F} = (F^{(1)} \otimes \dots \otimes F^{(N)})$ the linear transformation from $\mathbb{R}^{M_1 \otimes \dots \otimes M_N}$ to $\mathbb{R}^{M'_1 \otimes \dots \otimes M'_N}$, such that for any elementary tensor \mathcal{A} , with notations as above, it holds that:

$$\mathbf{F}(\mathcal{A}) = F^{(1)}(\mathbf{v}^{(1)}) \otimes \dots \otimes F^{(N)}(\mathbf{v}^{(N)})$$

$\mathbf{F}(\mathcal{A})$ is defined for a general tensor \mathcal{A} through its rank-1 decomposition comprising elementary tensors and applying \mathbf{F} on each of them, which can be shown to be equivalent to

$$\mathbf{F}(\mathcal{A})_{k_1, \dots, k_N} = \sum_{d_1=1}^{M_1} \dots \sum_{d_N=1}^{M_N} \mathcal{A}_{d_1, \dots, d_N} \prod_{i=1}^N F_{k_i, d_i}^{(i)} \quad (7)$$

A central concept in tensor analysis is that of *tensor matricization*. Let $P \sqcup Q = [N]$ be a disjoint partition of its indices, such that $P = \{p_1, \dots, p_{|P|}\}$ with $p_1 < \dots < p_{|P|}$, and $Q = \{q_1, \dots, q_{|Q|}\}$ with $q_1 < \dots < q_{|Q|}$. The matricization of \mathcal{A} with respect to the partition $P \sqcup Q$, denoted by $\llbracket \mathcal{A} \rrbracket_{P, Q}$, is the $\left(\prod_{t=1}^{|P|} M_{p_t} \right)$ -by- $\left(\prod_{t=1}^{|Q|} M_{q_t} \right)$ matrix holding the entries of \mathcal{A} , such that for all $i \in [N]$ and $d_i \in [M_i]$ the entry $\mathcal{A}_{d_1, \dots, d_N}$ is placed in row index $1 + \sum_{t=1}^{|P|} (d_{p_t} - 1) \prod_{t'=t+1}^{|P|} M_{p_{t'}}$ and column index $1 + \sum_{t=1}^{|Q|} (d_{q_t} - 1) \prod_{t'=t+1}^{|Q|} M_{q_{t'}}$. Applying the matricization operator $\llbracket \cdot \rrbracket_{P, Q}$ on the tensor product operator results in the *Kronecker Product*, i.e. for an N -ordered tensor \mathcal{A} , a K -ordered tensor \mathcal{B} , and the partition $P \sqcup Q = [N + K]$, it holds that

$$\llbracket \mathcal{A} \otimes \mathcal{B} \rrbracket_{P, Q} = \llbracket \mathcal{A} \rrbracket_{P \cap [N], Q \cap [N]} \odot \llbracket \mathcal{B} \rrbracket_{(P-N) \cap [K], (Q-N) \cap [K]}$$

where $P - N$ and $Q - N$ are simply the sets obtained by subtracting the number N from every element of P or Q , respectively. In concrete terms, the Kronecker product for the matrices $A \in \mathbb{R}^{M_1 \times M_2}$ and $B \in \mathbb{R}^{N_1 \times N_2}$ results in the matrix $A \odot B \in \mathbb{R}^{M_1 N_1 \times M_2 N_2}$ holding $A_{ij} B_{kl}$ in row index $(i-1)N_1 + k$ and column index $(j-1)N_2 + l$. An important property of the Kronecker product is that $\text{rank}(A \odot B) = \text{rank}(A) \cdot \text{rank}(B)$. Typically, when wish to compute $\text{rank}(\llbracket \mathcal{A} \rrbracket_{P, Q})$, we will first decompose it to a Kronecker product of matrices.

For a linear transform \mathbf{F} , as defined above in eq.7, and a partition $P \sqcup Q$, if $F^{(1)}, \dots, F^{(N)}$ are non-singular matrices, then \mathbf{F} is invertible and the matrix rank of $\llbracket \mathcal{A} \rrbracket_{P, Q}$ equals to the matrix rank of $\llbracket \mathbf{F}(\mathcal{A}) \rrbracket_{P, Q}$ (see proof in Hackbusch (2012)). Finally, we define the concept of *grid tensors*: for a function $f: \mathbb{R}^s \times \dots \times \mathbb{R}^s \rightarrow \mathbb{R}$ and a set of *template vectors* $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(M)} \in \mathbb{R}^s$, the N -order grid tensor $\mathcal{A}(f)$ is defined by $(\mathcal{A}(f))_{d_1, \dots, d_N} = f(\mathbf{x}^{(d_1)}, \dots, \mathbf{x}^{(d_N)})$.

In the context of ConvACs, circuits and the functions they can realize are typically examined through the matricization of the grid tensors they induce. The following is a succinct summary of the relevant known results used in our proofs – for a more detailed discussion, see previous works (Cohen et al., 2016a; Cohen and Shashua, 2016; 2017). Using the same notations from eq. 1 describing a general ConvAC, let \mathcal{A}^y be the coefficients tensor of order N and dimension M in each mode, and let $f_{\theta_1}, \dots, f_{\theta_M}: \mathbb{R}^s \rightarrow \mathbb{R}$ be a set of M representation functions (see sec. 2). Under the above definitions, a non-overlapping ConvAC can be said to decompose the coefficients tensor \mathcal{A}^y . Different network architectures correspond to known tensor

decompositions: shallow networks corresponds to rank-1 decompositions, and deep networks corresponds to Hierarchical Tucker decompositions.

In Cohen and Shashua (2017), it was found that the matrix rank of the matricization of the coefficients tensors \mathcal{A}^y could serve as a bound for the size of networks decomposing \mathcal{A}^y . For the conventional non-overlapping ConvAC and the contiguous “low-high” partition $P = \{1, \dots, N/2\}, Q = \{N/2 + 1, \dots, N\}$ of $[N]$, the rank of the matricization $\llbracket \mathcal{A}^y \rrbracket_{P,Q}$ serves as a lower-bound on the number of channels of the next to last layer of any network which decomposes the coefficients tensor \mathcal{A} . In the common case of square inputs, i.e. the input is of shape $W \times W$ and $N = W^2$, it is more natural to represent indices by pairs (j, i) denoting the spatial location of each “patch” $\mathbf{x}_{(j,i)}$, where the first argument denotes the vertical location and the second denotes the horizontal location. Under such setting the equivalent “low-high” partitions are either the “left-right” partition, i.e. $P = \{(j, i) | j \leq \frac{H}{2}\}, Q = \{(j, i) | j > \frac{H}{2}\}$, or the “top-bottom” partition, i.e. $P = \{(j, i) | i \leq \frac{W}{2}\}, Q = \{(j, i) | i > \frac{W}{2}\}$. More generally, when considering networks using other *pooling geometries*, i.e. not strictly contiguous pooling windows, then for each pooling geometry there exists a corresponding partition $P \cup Q$ such that $\text{rank} \llbracket \mathcal{A}^y \rrbracket_{P,Q}$ serves as its respective lower-bound.

Though the results in Cohen and Shashua (2017) are strictly based on the matricization rank of the coefficients tensors, they can be transferred to the matricization rank of grid tensors as well. Grid tensors were first considered for analyzing ConvACs in Cohen and Shashua (2016). For a set of M template vectors $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(M)} \in \mathbb{R}^s$, we define the matrix $F \in \mathbb{R}^{M \times M}$ by $F_{ij} = f_{\theta_j}(\mathbf{x}^{(i)})$. With the above notations in place, we can write the grid tensor $\mathcal{A}(\mathbf{h}_y)$ for the function $\mathbf{h}_y(\mathbf{x}_1, \dots, \mathbf{x}_N)$ as:

$$\begin{aligned} \mathcal{A}(\mathbf{h}_y)_{k_1, \dots, k_N} &= \mathbf{h}_y(\mathbf{x}^{(k_1)}, \dots, \mathbf{x}^{(k_N)}) \\ &= \sum_{d_1, \dots, d_N=1}^M \mathcal{A}_{d_1, \dots, d_N}^y \prod_{i=1}^N f_{\theta_{d_i}}(\mathbf{x}^{(k_i)}) \\ &= \sum_{d_1, \dots, d_N=1}^M \mathcal{A}_{d_1, \dots, d_N}^y \prod_{i=1}^N F_{k_i, d_i} \\ &\Rightarrow \mathcal{A}(\mathbf{h}_y) = (F \otimes \dots \otimes F)(\mathcal{A}^y) \end{aligned}$$

If the representation functions are linearly independent and continuous, then we can choose the template vectors such that F is non-singular (see Cohen and Shashua (2016)), which according to the previous discussion on tensor matricization, means that for any partition $P \cup Q$ and any coefficients tensor \mathcal{A}^y , it holds that $\text{rank}(\llbracket \mathcal{A}^y \rrbracket_{P,Q}) = \text{rank}(\llbracket \mathcal{A}(\mathbf{h}_y) \rrbracket_{P,Q})$. Thus, any lower bound on the matricization rank of the grid tensor translates to a lower bound on the matricization rank of the coefficients tensors, which in turn serves as lower bound on the size of non-overlapping ConvACs. We can summarize the above detailed discussion with the following two lemmas:

Lemma 1. Let $\mathbf{h}_y : \mathbb{R}^s \times \dots \times \mathbb{R}^s \rightarrow \mathbb{R}$ be a score function. Fix a set of M linearly independent and continuous representation functions, and choose template vectors such that F is non-singular. Let $f_y : \mathbb{R}^s \times \dots \times \mathbb{R}^s \rightarrow \mathbb{R}$ be the function realized by a non-overlapping ConvAC with standard pooling geometry, such that the induced grid tensors hold $\mathcal{A}(f_y) = \mathcal{A}(\mathbf{h}_y)$. If $P \cup Q$ is either the “left-right” or “top-bottom” partition, as described above, then the number of channels of its next to last layer must be greater than or equal to $\text{rank}(\llbracket \mathcal{A}(\mathbf{h}_y) \rrbracket_{P,Q})$.

Lemma 2. Under the same conditions as lemma 1, if for all partitions $P \cup Q$ such that $|P| = |Q| = N/2$ it holds that $\text{rank}(\llbracket \mathcal{A}(\mathbf{h}_y) \rrbracket_{P,Q}) \geq T$, then any ConvAC with any pooling geometry must have at least T channels in its next to last layer.

Proof. (of lemma 1 and lemma 2) For the proofs of the base results with respect to the coefficients tensor, see Cohen and Shashua (2017). To prove it is possible to choose the template vectors such that F is non-singular, see Cohen and Shashua (2016). To prove that if F is non-singular, then the grid tensor and the coefficients tensor have the same matricization rank, see lemma 5.6 in Hackbusch (2012). \square

Finally, we quote the following lemma regarding the prevalence of the maximal matrix rank for matrices whose entries are polynomial functions:

Lemma 3. Let $M, N, K \in \mathbb{N}$, $1 \leq r \leq \min\{M, N\}$ and a polynomial mapping $A : \mathbb{R}^K \rightarrow \mathbb{R}^{M \times N}$, i.e. for every $i \in [M]$ and $j \in [N]$ it holds that $A_{ij} : \mathbb{R}^K \rightarrow \mathbb{R}$ is a polynomial function. If there exists a point $\mathbf{x} \in \mathbb{R}^K$ such that $\text{rank}(A(\mathbf{x})) \geq r$, then the set $\{\mathbf{x} \in \mathbb{R}^K | \text{rank} A(\mathbf{x}) < r\}$ has zero measure (with respect to the Lebesgue measure over \mathbb{R}^K).

Proof. See Sharir et al. (2016). \square

With the above definitions and lemmas in place, we are ready to prove the propositions and theorems from the body of the article.

A.2 PROOF OF PROPOSITION 1

Proposition 1 is a direct corollary of the following two claims:

Claim 1. Let $f: \mathbb{R}^{D \times W \times W} \rightarrow \mathbb{R}^{C \times N \times N}$ be a function realized by a single GC layer with $R \times R$ local receptive field, $S \times S$ stride, and C output channels, that is parameterized by $\{(\mathbf{w}^{(c)}, \mathbf{b}^{(c)})\}_{c=1}^C$. For all $\tilde{R} \geq R$, a GC layer with $\tilde{R} \times \tilde{R}$ local receptive field, $S \times S$ stride, and C output channels, parameterized by $\{(\tilde{\mathbf{w}}^{(c)}, \tilde{\mathbf{b}}^{(c)})\}_{c=1}^C$, could also realize f . The same is true for the unshared case of both layers.

Proof. The claim is trivially satisfied by setting $\tilde{\mathbf{w}}^{(c)}$ such that it is equal to $\mathbf{w}^{(c)}$ in all matching coordinates, while using zeros for all other coordinates. Similarly, we set $\tilde{\mathbf{b}}^{(c)}$ to be equal to $\mathbf{b}^{(c)}$ in all matching coordinates, while using ones for all other coordinates. \square

Claim 2. Let $f: \mathbb{R}^{D \times W \times W} \rightarrow \mathbb{R}^{D \times N \times N}$ be a function realized by a GC layer with $R \times R$ local receptive field and 1×1 stride, parameterized by $\{(\mathbf{w}^{(c)}, \mathbf{b}^{(c)})\}_{c=1}^C$. There there exist an assignment to (\mathbf{w}, \mathbf{b}) such that f is the identity function $f(X) = X$. The same is true for the unshared case of both layers.

Proof. From claim 1 it is sufficient to show the above holds for $R = 1$. Indeed, setting

$$w_d^{(c)} = 1_{[d=c]} = \begin{cases} 1 & d = c \\ 0 & d \neq c \end{cases}$$

and $\mathbf{b}^{(c)} \equiv \mathbf{0}$ satisfies the claim. \square

A.3 PROOF OF THEOREM 1

Given the structure of an overlapping network, we wish to show that with probability 1, we obtain a function that cannot be realized by a non-overlapping network, unless the number of channels of its next to last layer is above a certain threshold. As discussed in sec. 3.2, instead of finding a lower bound for realizing a function, we find a lower bound for inducing the same grid tensor. According to lemma 1, the matricization rank according to either the “left-right” or “top-bottom” matricizations (see definitions in sec. A.1) of a grid tensor serves as a lower bound on the size of non-overlapping networks. Thus, it is enough to show that the matricization rank of grid tensors induced by an overlapping network of a given structure are greater than the bound given by theorem 1. In order for us to prove that when we randomize the parameters of the network, the matricization rank holds the above inequality with probability 1, is equivalent to proving that the set of parameters, for which this is not the case, is of measure zero, with respect to the Lebesgue measure over the Euclidean space of the parameters. Finally, the entries of the grid tensor induced by an overlapping network are polynomial functions of its parameters, and according to lemma 3, if we find just a single example for which the matricization rank is greater than the desired lower bound, then this is true almost everywhere, i.e. with probability 1. Giving the last remark, our proof is simply the construction of such an example.

We begin by fixing a set of M continuous and linearly independent representation functions, $f_{\theta_1}, \dots, f_{\theta_M}: \mathbb{R}^s \rightarrow \mathbb{R}$, and choose template vectors $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(M)} \in \mathbb{R}^s$ according to the discussion in sec. A.1, such that F is a non-singular matrix defined by $F_{ij} = f_{\theta_j}(\mathbf{x}^{(i)})$. We additionally assume in the following claims that the output of the representation layer is of width and height equal to $W \in \mathbb{N}$, where W is an even number – the claims and proofs can however be easily adapted to the more general case.

Let Φ be a ConvAC with representation layer defined according to above, followed by L GC layers, where the l ’th layer has a local receptive field of $R^{(l)} \times R^{(l)}$, a stride of $S^{(l)} \times S^{(l)}$, and $D^{(l)}$ output channels. We first construct our example, that achieves the desired matricization rank, for the simpler case where the first layer following the representation layer has a local receptive field large enough, i.e. when it is larger than $\frac{W}{2}$. Recall that for the first layer the total receptive field is equal to its local receptive field. In the context of theorem 1, this first layer satisfies the conditions necessary to produce the lower bound given in the theorem.

The specific construction is presented in the following claim, which relies on utilizing the large local receptive field to match each spatial location in the left side of the input with one from the right side, such that for each such pair, the respective output of the first layer will represent a mostly diagonal matrix. We then set the rest of the parameters such that the output of the entire network is defined by a tensor product of mostly diagonal matrices. Since the matricization rank of the tensor product of matrices is equal to the product of the individual ranks, it results in an exponential form of the rank as is given in the theorem.

Claim 3. Let Φ be a ConvAC as defined above, ending with a single scalar output. For all $l \in [L]$, the parameters of the l -th GC layer are denoted by $\{(\mathbf{w}^{(l,c)} \in \mathbb{R}^{D^{(l-1)} \times R^{(l)} \times R^{(l)}}, \mathbf{b}^{(l,c)} \in \mathbb{R}^{R^{(l)} \times R^{(l)}})\}_{c \in [D^{(l)}]}$. Let $\mathbf{h}(\mathbf{x}_1, \dots, \mathbf{x}_N)$ be the function realized by Φ . Additionally define $R \equiv R^{(1)}$, $S \equiv S^{(1)}$ and $D \equiv$

$\min\{D^{(1)}, M\}$. If $R > \frac{W}{2}$, and the weights $\mathbf{w}^{(1,c)}$ and biases $\mathbf{b}^{(1,c)}$ of the first GC layer are set to:

$$w_{mji}^{(1,c)} = \begin{cases} -\alpha (F^{-1})_{m,c} & c \leq D \text{ and } (j, i) \in \{(1, 1), (\rho, \tau)\} \\ 0 & \text{Otherwise} \end{cases}$$

$$b_{ji}^{(1,c)} = \begin{cases} \beta & c \leq D \text{ and } (j, i) \in \{(1, 1), (\rho, \tau)\} \\ 1 & c \leq D \text{ and } (j, i) \notin \{(1, 1), (\rho, \tau)\} \\ 0 & \text{Otherwise} \end{cases}$$

where $\beta = \frac{2\alpha}{D}$, then there exists an assignment to α and the parameters of the other GC layers such that $\text{rank}(\llbracket \mathcal{A}(\mathbf{h}) \rrbracket_{P,Q}) = D \lfloor \frac{W-R}{S} + 1 \rfloor \cdot \lceil \frac{W}{S} \rceil$, where $P \cup Q$ is either the “left-right” or “top-bottom” partition, and (ρ, τ) equals to $(1, R)$ or $(R, 1)$, respectively.

Proof. The proof for either the “left-right” or “top-bottom” partition is completely symmetric, thus it is enough to prove the claim for the “left-right” case, where $(\rho, \tau) = (1, R)$. We wish to compute the entry $\mathcal{A}(\mathbf{h})_{d_{(1,1)}, \dots, d_{(W,W)}}$ of the induced grid tensor for arbitrary indices $d_{(1,1)}, \dots, d_{(W,W)}$. Let $O \in \mathbb{R}^{M \times W \times W}$ be the 3-order tensor output of the representation layer, where $O_{m,j,i} = F_{d_{(j,i)},m}$ for the aforementioned indices and for all $1 \leq i, j \leq W$ and $m \in [M]$.

We begin by setting the parameters of all layers following the first GC layer, such that they are equal to computing the sum along the channels axis of the output of the first GC layer, followed by a global product of all of the resulting sums. To achieve this, we can first assume w.l.o.g. that these layers are non-overlapping through proposition 1. We then set the parameters of the second GC layer to $\mathbf{w}^{(2,c)} = \mathbf{1}$ and $\mathbf{b}^{(2,c)} \equiv \mathbf{0}$, i.e. all ones and all zeros, respectively, which is equivalent to taking the sum along the channels axis for each spatial location, followed by taking the products over non-overlapping local receptive fields of size $R^{(2)}$. For the other layers, we simply set them as to take just the output of the first channel of the output of the preceding layer,

which is equal to setting their parameters to $w_{dji}^{(l,c)} = \begin{cases} 1 & d = 1 \\ 0 & d \neq 1 \end{cases}$ and $\mathbf{b}^{(l,c)} \equiv \mathbf{0}$. Setting the parameters as described above results in:

$$(\mathcal{A}(\mathbf{h}))_{d_{(1,1)}, \dots, d_{(W,W)}} = \prod_{\substack{0 \leq uS < W \\ 0 \leq vS < W}} \sum_{c=1}^{D^{(1)}} \prod_{j,i=1}^R \underbrace{\left(b_{ji}^{(1,c)} + \sum_{m=1}^M w_{mji}^{(1,c)} O_{m,uS+j,vS+i} \right)}_{g(u,v,c,j,i)} \quad (8)$$

where we extended O with zero-padding for the cases where $uS + j > W$ or $vS + i > W$, as mentioned in sec. 3.1. Next, we go through the technical process of reducing eq. 8 to a product of matrices.

Substituting the values of $w_{dji}^{(1,c)}$ and $b_{ji}^{(1,c)}$ with those defined in the claim, and computing the value of $g(u, v, c, j, i)$ results in:

$$g(u, v, c, j, i) = \begin{cases} \beta - \alpha \sum_{m=1}^M (F^{-1})_{m,c} F_{d_{(uS+j,vS+i)},m} & c \leq D \text{ and } vS + R \leq W \text{ and } (j, i) \in \{(1, 1), (1, R)\} \\ \beta - \alpha \sum_{m=1}^M (F^{-1})_{m,c} F_{d_{(uS+j,vS+i)},m} & c \leq D \text{ and } vS + R > W \text{ and } (j, i) = (1, 1) \\ \beta & c \leq D \text{ and } vS + R > W \text{ and } (j, i) = (1, R) \\ 1 & c \leq D \text{ and } (j, i) \notin \{(1, 1), (1, R)\} \\ 0 & c > D \end{cases}$$

$$= \begin{cases} \beta - \alpha (FF^{-1})_{d_{(uS+1,vS+i)},c} & c \leq D \text{ and } vS + R \leq D \text{ and } (j, i) \in \{(1, 1), (1, R)\} \\ \beta - \alpha (FF^{-1})_{d_{(uS+1,vS+i)},c} & c \leq D \text{ and } vS + R > W \text{ and } (j, i) = (1, 1) \\ \beta & c \leq D \text{ and } vS + R > W \text{ and } (j, i) = (1, R) \\ 1 & c \leq D \text{ and } (j, i) \notin \{(1, 1), (1, R)\} \\ 0 & c > D \end{cases}$$

from which we derive:

$$f(u, v) \equiv \sum_{c=1}^{D^{(1)}} \prod_{j,i=1}^R g(u, v, c, j, i) = \begin{cases} D\beta^2 - \alpha\beta(1_{[d_{(uS+1,vS+1)} \leq D]} + 1_{[d_{(uS+1,vS+R)} \leq D]}) & vS + R \leq W \\ + \alpha^2 1_{[d_{(uS+1,vS+1)} = d_{(uS+1,vS+R)} \leq D]} & \\ D\beta^2 - \alpha\beta 1_{[d_{(uS+1,vS+1)} \leq D]} & vS + R > W \end{cases}$$

where $(\mathcal{A}(\mathbf{h}))_{d_{(1,1)}, \dots, d_{(W,W)}} = \prod_{\substack{0 \leq uS < W \\ 0 \leq vS < W}} f(u, v)$.

At this point we branch into two cases. If S divides $R - 1$, then for all $u, v \in \mathbb{N}$ such that $vS + R \leq W$ and $uS < W$, the above expression for $f(u, v)$ and $f(u, v + \frac{R-1}{S})$ depends only on the indices $d_{(uS+1, vS+1)}$ and $d_{(uS+1, vS+R)}$, while these two indices affect only the aforementioned expressions. By denoting $A_{d_{(uS+1, vS+1)}, d_{(uS+1, vS+R)}}^{(u, v)} = f(u, v) \cdot f(u, v + \frac{R-1}{S})$, we can write it as:

$$A_{ij}^{(u, v)} = \begin{cases} (D\beta^2 - 2\alpha\beta + \alpha^2 1_{[i=j]}) (D\beta^2 - \alpha\beta) & i, j \leq D \\ (D\beta^2 - \alpha\beta) (D\beta^2) & i \leq D \text{ and } j > D \\ (D\beta^2 - \alpha\beta)^2 & i > D \text{ and } j \leq D \\ (D\beta^2)^2 & i, j > D \end{cases}$$

where $i, j \in [M]$ stand for the possible values of $d_{(uS+1, vS+1)}$ and $d_{(uS+1, vS+R)}$, respectively. Substituting $\beta = \frac{2\alpha}{D}$, as stated in the claim, and setting $\alpha = (\frac{D}{2})^{1/4}$, results in:

$$A_{ij}^{(u, v)} = \begin{cases} 1_{[i=j]} & i, j \leq D \\ \frac{4}{D} & i \leq D \text{ and } j > D \\ \frac{2}{D} & i > D \text{ and } j \leq D \\ \frac{8}{D} & i, j > D \end{cases}$$

which means $\text{rank}(A^{(u, v)}) = D$. Since $(\mathcal{A}(\mathbf{h}))_{d_{(1,1)}, \dots, d_{(W, W)}}$ equals to the product

$\prod_{\substack{0 \leq uS < W \\ 0 \leq vS \leq W-R}} A_{d_{(uS+1, vS+1)}, d_{(uS+1, vS+R)}}^{(u, v)}$, then $\llbracket \mathcal{A}(\mathbf{h}) \rrbracket_{P, Q}$ equals to the Kronecker product of the matrices in $\{A^{(u, v)} | 0 \leq uS < W, 0 \leq vS \leq W-R\}$, up to permutation of its rows and columns, which do not affect its matrix rank. Thus, the matricization rank of $\mathcal{A}(\mathbf{h})$ satisfies:

$$\text{rank}(\llbracket \mathcal{A}(\mathbf{h}) \rrbracket_{P, Q}) = D |\{A^{(u, v)} | 0 \leq uS < W, 0 \leq vS \leq W-R\}| = D \lfloor \frac{W-R}{S} + 1 \rfloor \cdot \lceil \frac{W}{S} \rceil$$

which proves the claim for this case.

If S does not divide $R - 1$, then for all $u, v \in \mathbb{N}$, such that $vS + R \leq W$ and $uS < W$, it holds that $f(u, v)$ depends only on the indices $d_{(uS+1, vS+1)}$ and $d_{(uS+1, vS+R)}$, and they affect only $f(u, v)$. Additionally, for all $u, v \in \mathbb{N}$, such that $W < vS + R$, $vS < W$ and $uS < W$, it holds that $f(u, v)$ depends only on the index $d_{(uS+1, vS+1)}$, and this index affects only $f(u, v)$. Let us denote $\mathcal{A}_{d_{(uS+1, vS+1)}, d_{(uS+1, vS+R)}}^{(u, v)} = f(u, v)$ for $vS + R \leq W$:

$$A_{ij}^{(u, v)} = \begin{cases} (D\beta^2 - 2\alpha\beta + \alpha^2 1_{[i=j]}) & i, j \leq D \\ (D\beta^2 - \alpha\beta) & i \leq D \text{ and } j > D \\ (D\beta^2 - \alpha\beta) & i > D \text{ and } j \leq D \\ (D\beta^2) & i, j > D \end{cases}$$

which by setting $\beta = \frac{2\alpha}{D}$ and $\alpha = 1$ results in:

$$A_{ij}^{(u, v)} = \begin{cases} 1_{[i=j]} & i, j \leq D \\ \frac{2}{D} & i \leq D \text{ and } j > D \\ \frac{2}{D} & i > D \text{ and } j \leq D \\ \frac{4}{D} & i, j > D \end{cases}$$

which means $\text{rank}(A^{(u, v)}) = D$. For $vS + R > W$, and we can define the vector $\mathbf{a}_{d_{(uS+1, vS+1)}}^{(u, v)} = f(u, v)$, which by using the same values of α and β results in:

$$\mathbf{a}_i^{(u, v)} = \begin{cases} \frac{2}{D} & i \leq D \\ \frac{4}{D} & i > D \end{cases}$$

By viewing $\mathbf{a}^{(u, v)}$ as either a column or row vector, depending on whether $d_{(uS+1, vS+1)} \in P$ or $d_{(uS+1, vS+1)} \in Q$, respectively, it holds that $\llbracket \mathcal{A}(\mathbf{h}) \rrbracket_{P, Q}$ equals to the Kronecker product of the matrices in $\{A^{(u, v)}\}_{\substack{0 \leq uS < W \\ 0 \leq vS \leq W-R}} \cup \{\mathbf{a}^{(u, v)}\}_{\substack{0 \leq uS < W \\ W-R < vS < W}}$, up to permutations of its rows and columns, which do not affect the rank. Since $\mathbf{a}^{(u, v)} \neq \mathbf{0}$ then $\text{rank}(\mathbf{a}^{(u, v)}) = 1$, which means the matricization rank of $\mathcal{A}(\mathbf{h})$ once again holds:

$$\text{rank}(\llbracket \mathcal{A}(\mathbf{h}) \rrbracket_{P, Q}) = D |\{A^{(u, v)} | 0 \leq uS < W, 0 \leq vS \leq W-R\}| = D \lfloor \frac{W-R}{S} + 1 \rfloor \cdot \lceil \frac{W}{S} \rceil$$

□

In the preceding claim we have describe an example for the case where the total receptive of the first GC layer is already large enough for satisfying the conditions of the theorem. In the following claim we extend this result for the general case. This is accomplished by showing that a network comprised of just L GC layers with local receptive fields $\{R^{(l)} \times R^{(l)}\}_{l \in [L]}$, strides $\{S^{(l)} \times S^{(l)}\}_{l \in [L]}$, and output channels $\{D^{(l)}\}_{l \in [L]}$, can effectively compute the same output as the first GC layer from claim 3, for all inputs³.

Recall that the layer from claim 3 performs an identical transformation on each $M \times 1 \times 1$ patch from its input, followed by taking the point-wise product of far-away pairs of transformed patches. Thus, the motivation behind the specific construction we use, is to use the first of the L layers to perform this transformation, while using half of its output channels for storing the transformed patch from the same location, and the other half for storing a transformed patch, but from a location farthest to the right, constrained by its local receptive field. This is equal to having one set of transformed patches sitting still, while another “shifted” set of transformed patches. The other layers simply pass the the first half of the channels as is, using an identity operation as defined in claim 2, while continuously shifting the other half of the channels more and more to the left, bringing faraway patches closer together. Finally, at the last layer we take both halves and multiple them together.

Claim 4. *Let Φ be a ConvAC comprised of just L GC layers as described above. The output of Φ is not limited to a scalar value. Assume the total stride of the L -th GC layer is greater than $W/2$, and let $T_S^{(L)}$ and $T_R^{(L, \alpha)}$ be the total stride and the α -minimal total receptive field, respectively, for $\alpha = \lfloor W/2 \rfloor + 1$. Let Ψ be a ConvAC comprised of a single GC layer with local receptive field $R \equiv T_R^{(L, \alpha)}$, stride $S \equiv T_S^{(L)}$, output channels $D \equiv \min\{\frac{1}{2} \min_{1 \leq l < L} D^{(l)}, D^{(L)}\}$, where its weights and biases are set to the following:*

$$w_{mji}^{(c)} = \begin{cases} A_{m,c} & (j, i) \in \{(1, 1), (\rho, \tau)\} \\ 0 & \text{Otherwise} \end{cases}$$

$$b_{ji}^{(c)} = \begin{cases} \beta_c & (j, i) \in \{(1, 1), (\rho, \tau)\} \\ 1 & \text{Otherwise} \end{cases}$$

for $\beta \in \mathbb{R}^D$, $A \in \mathbb{R}^{M \times D}$ and $(\rho, \tau) \in \{(1, R), (R, 1)\}$. Then, there exists a set of weights to the layers of Φ such that for every input X , the output of Φ is equivalent to the output of Ψ for channels $\leq D$, and zero otherwise.

Proof. The two possible cases for (ρ, τ) are completely symmetric, thus it is enough to prove the claim just for $(\rho, \tau) = (1, R)$. Additionally, we can assume w.l.o.g. that $\forall l, R^{(l)} > 1$, by setting any 1×1 layer to act as pass-through according to claim 2, and also assume that the α -minimal total receptive field is exactly equal to the total receptive field of the L -th layer, by applying claim 1 to realize an equivalent network with smaller windows. Finally, the case for $L = 1$ is trivial, and thus we assume $L > 1$.

Let us set the parameters $\{\mathbf{w}^{(l,k)}, \mathbf{b}^{(l,k)}\}$ of the layers of Φ as follows:

$$w_{dji}^{(l,k)} = \begin{cases} -A_{d,k} & (l = 1) \text{ and } (1 \leq k \leq D) \text{ and } (j, i) = (1, 1) \\ -A_{d,k-D} & (l = 1) \text{ and } (D < k \leq 2D) \text{ and } (j, i) = (1, R^{(l)}) \\ 1_{[d=k]} & (1 < l < L) \text{ and } (1 \leq k \leq D) \text{ and } (j, i) = (1, 1) \\ 1_{[d=k]} & (1 < l < L) \text{ and } (D < k \leq 2D) \text{ and } (j, i) = (1, R^{(l)}) \\ 1_{[d=k]} & (l = L) \text{ and } (1 \leq k \leq D) \text{ and } (j, i) = (1, 1) \\ 1_{[d=M+k]} & (l = L) \text{ and } (1 \leq k \leq D) \text{ and } (j, i) = (1, R^{(l)}) \\ 0 & \text{Otherwise} \end{cases}$$

$$b_{ji}^{(l,k)} = \begin{cases} \beta_k & (l = 1) \text{ and } (1 \leq k \leq D) \text{ and } (j, i) = (1, 1) \\ \beta_{k-D} & (l = 1) \text{ and } (D < k \leq 2D) \text{ and } (j, i) = (1, R^{(l)}) \\ 0 & (1 < l < L) \text{ and } (1 \leq k \leq D) \text{ and } (j, i) = (1, 1) \\ 0 & (1 < l < L) \text{ and } (D < k \leq 2D) \text{ and } (j, i) = (1, R^{(l)}) \\ 0 & (l = L) \text{ and } (1 \leq k \leq D) \text{ and } (j, i) = (1, 1) \\ 0 & (l = L) \text{ and } (1 \leq k \leq D) \text{ and } (j, i) = (1, R^{(l)}) \\ 1 & \text{Otherwise} \end{cases}$$

It is left to prove the above satisfies the claim. Let $O^{(l)} \in \mathbb{R}^{D^{(l)} \times W^{(l)} \times W^{(l)}}$ be the output of the l -th layer, for $l \in [0, \dots, L]$, where $W^{(l)}$ is the width and height of the output. We additionally assume that for indices beyond the bounds of $[D^{(l)}] \times [W^{(l)}] \times [W^{(l)}]$ the value of $O^{(l)}$ is zero, i.e. we assume zero padding when applying the convolutional operation of the GC layer. We extend the definition for $l = 0$, by setting $D^{(0)} \equiv M$

³Notice that in this context, there is no representation layer, and the input can be any 3-order tensor.

and $W^{(0)} \equiv W$, where we identify $O^{(0)}$ with the input to the network Φ . Given the above, the output of the first layer for $k \in [D^{(1)}]$ and $0 \leq u, v < W^{(1)}$, is as follows:

$$\begin{aligned} O_{k,u+1,v+1}^{(1)} &= \prod_{j,i=1}^{R^{(1)}} \left(b_{ji}^{(1,k)} + \sum_{d=1}^{D^{(0)}} w_{dji}^{(1,k)} O_{d,uS_h+j,vS_w+i}^{(0)} \right) \\ &= \begin{cases} \beta_k + \sum_{d=1}^M A_{d,k} \cdot O_{d,uS_h+1,vS_w+1}^{(0)} & 1 \leq k \leq D \\ \beta_{k-D} + \sum_{d=1}^M A_{d,k-D} \cdot O_{d,uS_h+1,vS_w+R^{(1)}}^{(0)} & D < k \leq 2D \\ 1 & \text{Otherwise} \end{cases} \end{aligned}$$

We will show by induction that for $1 < l < L$, $k \in [D^{(l)}]$ and $0 \leq u, v < W^{(l)}$ the output of the l -th layer $O_{k,u+1,v+1}^{(l)}$ always equals to:

$$O_{k,u+1,v+1}^{(l)} = \begin{cases} O_{k,u\eta^{(l)}+1,v\eta^{(l)}+1}^{(1)} & k \leq D \\ O_{k,u\eta^{(l)}+1,v\eta^{(l)}+\xi^{(l)}}^{(1)} & D < k \leq 2D \\ 1 & \text{Otherwise} \end{cases}$$

where $\eta^{(l)} = \prod_{i=2}^l S^{(i)}$ and $\xi^{(l)} = R^{(l)} \cdot \eta^{(l-1)} + \sum_{k=2}^{l-1} (R^{(k)} - S^{(l)}) \cdot \eta^{(k-1)}$. It is trivial to verify that for $l = 2$ it indeed holds, since:

$$O_{k,u+1,v+1}^{(2)} = \begin{cases} O_{k,uS^{(2)}+1,vS^{(2)}+1}^{(1)} & k \leq D \\ O_{k,uS^{(2)}+1,vS^{(2)}+R^{(2)}}^{(1)} & D < k \leq 2D \\ 1 & \text{Otherwise} \end{cases}$$

where $\eta^{(2)} = S^{(2)}$ and $\xi^{(2)} = R^{(2)}$. Assume the claim holds up to $l-1$, and we will show it also holds for l :

$$\begin{aligned} O_{k,u+1,v+1}^{(l)} &= \prod_{j,i=1}^{R^{(l)}} \left(b_{ji}^{(l,k)} + \sum_{d=1}^{D^{(l-1)}} w_{dji}^{(l,k)} O_{d,uS^{(l)}+j,vS^{(l)}+i}^{(l-1)} \right) \\ &= \begin{cases} O_{k,uS^{(l)}+1,vS^{(l)}+1}^{(l-1)} & k \leq D \\ O_{k,uS^{(l)}+1,vS^{(l)}+R^{(l)}}^{(l-1)} & D < k \leq 2D \\ 1 & \text{Otherwise} \end{cases} \\ \text{Induction Hypothesis} \Rightarrow &= \begin{cases} O_{k,(uS^{(l)})\eta^{(l-1)}+1,(vS^{(l)})\eta^{(l-1)}+1}^{(1)} & k \leq D \\ O_{k,(uS^{(l)})\eta^{(l-1)}+1,(vS^{(l)}+R^{(l)}-1)\eta^{(l-1)}+\xi^{(l-1)}}^{(1)} & D < k \leq 2D \\ 1 & \text{Otherwise} \end{cases} \\ &= \begin{cases} O_{k,u\eta^{(l)}+1,v\eta^{(l)}+1}^{(1)} & k \leq D \\ O_{k,u\eta^{(l)}+1,v\eta^{(l)}+\xi^{(l)}}^{(1)} & D < k \leq 2D \\ 1 & \text{Otherwise} \end{cases} \end{aligned}$$

Were we used the fact that $\eta^{(l)} = S^{(l)}\eta^{(l-1)}$ and $\xi^{(l)} = R^{(l)}\eta^{(l-1)} + \xi^{(l-1)} - \eta^{(l-1)}$.

Finally, we show that $O_{k,u+1,v+1}^{(L)}$ for $k \leq D$ and $0 \leq u, v < W^{(L)}$ equals to the output of the single GC layer specified in the claim:

$$\begin{aligned}
O_{k,u+1,v+1}^{(L)} &= \prod_{j,i=1}^{R^{(L)}} \left(b_{ji}^{(L,k)} + \sum_{d=1}^{D^{(L-1)}} w_{dji}^{(L,k)} O_{d,uS^{(L)}+j,vS^{(L)}+i}^{(L-1)} \right) \\
&= O_{k,uS^{(L)}+1,vS^{(L)}+1}^{(L-1)} \cdot O_{k,uS^{(L)}+1,vS^{(L)}+R^{(L)}}^{(L-1)} \\
&= O_{k,u\eta^{(L)}+1,v\eta^{(L)}+1}^{(1)} \cdot O_{k,u\eta^{(L)}+1,v\eta^{(L)}+\xi^{(L)}}^{(1)} \\
&= \left(\beta_k + \sum_{d=1}^{D^{(0)}} A_{d,k} O_{d,u\eta^{(L)}S^{(1)}+1,v\eta^{(L)}S^{(1)}+1}^{(0)} \right) \\
&\quad \cdot \left(\beta_k + \sum_{d=1}^{D^{(0)}} A_{d,k} O_{d,u\eta^{(L)}S^{(1)}+1,(v\eta^{(L)}+\xi^{(L)}-1)S^{(1)}+R^{(1)}}^{(0)} \right) \\
&= \left(\beta_k + \sum_{d=1}^{D^{(0)}} A_{d,k} O_{d,uT_S^{(L)}+1,vT_S^{(L)}+1}^{(0)} \right) \left(\beta_k + \sum_{d=1}^{D^{(0)}} A_{d,k} O_{d,uT_S^{(L)}+1,vT_S^{(L)}+T_R^{(L)}}^{(0)} \right)
\end{aligned}$$

which is indeed equal to the single GC layer. For $k > D$, both the bias and the weights for the last layer are zero, and thus $O_{k,u+1,v+1}^{(L)} = 1$. \square

Finally, with the above two claims set in place, we can prove our main theorem:

Proof. (of theorem 1) Using claim 4 we can realize the networks from claim 3, for which the matricization rank for either partition equals to:

$$\overline{D} \left\lfloor \frac{W - T_R^{(\bar{L}, \lfloor W/2 \rfloor)}}{T_S} + 1 \right\rfloor \cdot \left\lceil \frac{W}{T_S} \right\rceil$$

Since for any matricization $[\mathcal{A}(\Psi)]_{P,Q}$ the entries of the matricization are polynomial functions with respect to the parameters of the network, then, according to lemma 3, the set of parameters of Φ , that does not attain the above rank, has zero measure. Since the union of zero measured sets is also of measure zero, then all parameters except a set of zero measure attain this matricization rank for both partitions at once. Finally, from lemma 1, for a network comprised solely of non-overlapping operations, the number of output channels of its next to last layer is lower-bounded by the matricization rank according to the “left-and-right” and “top-and-bottom” partitions, concluding the proof. \square

A.4 PROOF OF PROPOSITION 2

Following theorem 1, to compute the lower bound for the network described in proposition 2, we need to find the first layer for which its total receptive field is greater than $W/2$, and then estimate its total stride and its α -minimal total receptive field, for $\alpha = \lfloor W/2 \rfloor$. In the following claims we analyze the above properties of the given network:

Claim 5. *The total stride and total receptive field of the l -th $B \times B$ layer in the given network, i.e. the $(2l-1)$ -th GC layer after the representation layer, are given by the following equations:*

$$\begin{aligned}
T_S^{(2l-1)}(S^{(1)}, \dots, S^{(2l-1)}) &= 2^{l-1} \\
T_R^{(2l-1)}(R^{(1)}, S^{(1)}, \dots, R^{(2l-1)}, S^{(2l-1)}) &= (2B-1)2^{l-1} - B + 1
\end{aligned}$$

Proof. From eq. 3 it immediately follows that $T_S^{(2l-1)}(S^{(1)}, \dots, S^{(2l-1)}) = 2^{l-1}$. From eq. 4, the 2×2 stride 2 layers do not contribute to the receptive field as $R^{(2l)} - S^{(2l)} = 0$, which results in the following equation:

$$\begin{aligned}
T_R^{(2l-1)}(R^{(1)}, S^{(1)}, \dots, R^{(2l-1)}, S^{(2l-1)}) &= B \cdot 2^{l-1} + \sum_{i=1}^{l-1} (B-1)2^{i-1} \\
&= B \cdot 2^{l-1} + (B-1)(2^{l-1} - 1) \\
&= (2B-1)2^{l-1} - B + 1
\end{aligned}$$

\square

Claim 6. The α -minimal total receptive field for the l -th $B \times B$ layer in the given network, for $\alpha \in \mathbb{N}$ and $2^{l-1} \leq \alpha < 2^l - 1$, always equals $(\alpha + 1)$.

Proof. From eq. 5, the following holds:

$$\begin{aligned} T_R^{(2l-1, \alpha)} &= \underset{\substack{\forall i \in [l], 1 \leq t_{2i-1} \leq B \\ T_R^{(2l-1)}(t_1, 1, 2, 2, t_3, 1, \dots, t_{2l-1}, 1) > \alpha}}{\operatorname{argmin}} T_R^{(2l-1)}(t_1, 1, 2, 2, t_3, 1, \dots, t_{2l-1}, 1) \\ &= \underset{\substack{\forall i \in [l], 1 \leq t_{2i-1} \leq B \\ t_{2l-1} \cdot 2^{l-1} + \sum_{i=1}^{l-1} (t_{2i-1} - 1) 2^{i-1} > \alpha}}{\operatorname{argmin}} t_{2l-1} \cdot 2^{l-1} + \sum_{i=1}^{l-1} (t_{2i-1} - 1) 2^{i-1} \end{aligned}$$

Notice that the right term in the equation resembles a binary representation. If we limit t_{2i-1} to the set $\{1, 2\}$, this term can represent any number in the set $\{0, \dots, 2^{l-1} - 1\}$, and by choosing $t_{2l-1} = 1$, the complete term can represent any number in the set $\{2^{l-1}, \dots, 2^l - 1\}$, and specifically, for $2^{l-1} \leq \alpha < 2^l - 1$, there exists an assignment for $t_{2i-1} \in \{1, 2\}$ for $i \in [l-1]$ such that this terms equal $(\alpha + 1)$, and thus $T_R^{(2l-1, \alpha)} = \alpha + 1$. \square

With the above general properties for the given network, we can simplify the expression for the lower bound given in theorem 1:

Claim 7. If the l -th $B \times B$ layer in the given network satisfies $T_R^{(2l-1)}(R^{(1)}, S^{(1)}, \dots, R^{(2l-1)}, S^{(2l-1)}) > W/2$, then the lower bound given in theorem 1 equals to $M^{2^{2L-2l+1}}$

Proof. From the description of the network and the previous claims it holds that $\bar{D} = M$, $W = 2^L$, $T_S^{(2l-1)} = 2^{l-1}$, and $T_R^{(2l-1, \lfloor W/2 \rfloor)} = 2^{L-1} + 1$. Substituting all the above in eq. 6 results in:

$$\begin{aligned} (\text{Eq. 6}) &= M^{\left\lfloor \frac{2^L - 2^{L-1} - 1}{2^{l-1}} + 1 \right\rfloor \cdot \left\lceil \frac{2^L}{2^{l-1}} \right\rceil} \\ &= M^{\left\lfloor 2^{L-l} + 1 - \frac{1}{2^{l-1}} \right\rfloor \cdot 2^{L-l+1}} \\ &= M^{2^{2L-2l+1}} \end{aligned}$$

\square

With all of the above claims in place, we are ready to prove proposition 2:

Proof. (of proposition 2) From claim 5, we can infer which is the first $B \times B$ layer such that its receptive field is greater than $W/2$:

$$\begin{aligned} (2B - 1) \cdot 2^{l-1} - B + 1 &> 2^{L-1} \\ \Rightarrow l &> \log_2 \frac{2^L + 2B - 2}{2B - 1} \\ \Rightarrow l &= 1 + \left\lceil \log_2 \frac{2^L + 2B - 2}{2B - 1} \right\rceil \end{aligned}$$

Combining the above with claim 7, results in:

$$\begin{aligned} M^{2^{2L-2l+1}} &= M^{2^{2L-1-2 \left\lceil \log_2 \frac{2^L + 2B - 2}{2B - 1} \right\rceil}} \\ &\geq M^{2^{2L-1-2 \log_2 \frac{2^L + 2B - 2}{2B - 1}}} \\ (2^L \equiv W) \Rightarrow &= M^{\frac{W^2}{2} \left(1 + \frac{W-1}{2B-1}\right)^{-2}} \\ &= M^{\frac{(2B-1)^2}{2} \cdot \left(1 + \frac{2B-2}{W}\right)^{-2}} \end{aligned}$$

The limits and the special case for $B \leq \frac{W}{5} + 1$ are both direct corollaries of the above expressions. \square

A.5 PROOF OF THEOREM 2

The proof of theorem 2 follows the same general steps as the proof of theorem 1, where the concept of matricization rank, and specifically the “left-right” matricization, was used to find the desired lower bound according to lemma 1, and that based on lemma 3 it was enough to prove the existence of just a single case where the matricization rank was greater than the lower bound given in that theorem. For the proof of theorem 2, instead of lemma 1 we rely on lemma 2, which shows the minimal matricization rank, amongst all possible even matricizations of the induced grid tensor, results in a lower bound on the number of channels of the next to last hidden layer of a ConvAC realizing the same grid tensor, regardless of its pooling geometry.

We begin by proving an analogue of claim 3, where we show that for any given matricization of the grid tensor $\mathcal{A}(\mathbf{h})$, induced by the overlapping network realizing the function \mathbf{h} , the matricization rank is exponential. The motivation behind the construction, for when parameters are “unshared”, is that we can utilize the fact that there are separate sets of kernels, with local receptive fields the size of the input, for each spatial location. Thus, each kernel can “connect” the index (of the grid tensor) matching its spatial location, with almost any other index, and specifically such that the two indices come from different sets of the matricization $I \cup J$ of $\mathcal{A}(\mathbf{h})$. For the “shared” case, we simply use polynomially more output channels to simulate the “unshared” case.

Claim 8. *For an arbitrary even partition (I, J) of $\{(1, 1), \dots, (W, W)\}$, such that $|I| = |J| = \frac{W^2}{2}$, there exists an assignment to the parameters of the network given in theorem 2, for either the “unshared” or “shared” settings, such that $\text{rank}(\llbracket \mathcal{A}(\mathbf{h}) \rrbracket) = M^{\frac{W^2}{2}}$.*

Proof. Let (I, J) be an arbitrary even partition of $\{(1, 1), \dots, (W, W)\}$, such that $|I| = |J| = \frac{W^2}{2}$, $I = \{i_1, \dots, i_{|I|}\}$, and $J = \{j_1, \dots, j_{|J|}\}$, where for $k \in [\frac{W^2}{2}]$ it holds that $i_k < i_{k+1}$ and $j_k < j_{k+1}$ (using lexical ordering), and we assume w.l.o.g. that $i_1 = (1, 1)$. We define the set $\{(q_k, p_k)\}_{k=1}^{\frac{W^2}{2}}$ such that $q_k = i_k$ and $p_k = j_k$ if $i_k < j_k$, and otherwise $q_k = j_k$ and $p_k = i_k$.

We prove the “unshared” case first, where the parameters of the first GC layers are given by $\{(\mathbf{w}^{(c,u,v)}, \mathbf{b}^{(c,u,v)})\}_{c=1, u=1, v=1}^{D, W, W}$, for which we choose the following assignment:

$$w_{m,j,i}^{(c,u,v)} = \begin{cases} (F^{-1})_{m,c} & c \leq M \text{ and } \exists k \in [\frac{W^2}{2}], q_k = (u, v) \text{ and } (j, i) \in \{(1, 1), p_k - q_k + (1, 1)\} \\ 0 & \text{Otherwise} \end{cases}$$

$$b_{j,i}^{(c,u,v)} = \begin{cases} 0 & c \leq M \text{ and } \exists k \in [\frac{W^2}{2}], q_k = (u, v) \text{ and } (j, i) \in \{(1, 1), p_k - q_k + (1, 1)\} \\ 1_{[c \leq M]} & \text{Otherwise} \end{cases}$$

where for $u, v \in [W]$ such that $q_k = (u, v)$ it holds that $p_k - q_k + (1, 1) \in \{(1, 1), \dots, (W, W)\}$ because $q_k \leq p_k$. Similar to the proof of claim 3, we set all layers following the first GC layer such that the following equality holds:

$$(\mathcal{A}(\mathbf{h}))_{d_{(1,1)}, \dots, d_{(W,W)}} = \prod_{u,s=1}^W \sum_{c=1}^D \prod_{j,i=1}^W \left(b_{ji}^{(c,u,v)} + \sum_{m=1}^M w_{mji}^{(c,u,v)} O_{m,u+j-1,v+i-1} \right)$$

Which under the assignment to parameters we chose earlier, it results in:

$$\begin{aligned} (\mathcal{A}(\mathbf{h}))_{d_{(1,1)}, \dots, d_{(W,W)}} &= \prod_{k=1}^{W^2/2} \sum_{c=1}^D \left(\sum_{m=1}^M (F^{-1})_{m,c} F_{d_{q_k}, m} \right) \left(\sum_{m=1}^M (F^{-1})_{m,c} F_{d_{p_k}, m} \right) \\ &= \prod_{k=1}^{W^2/2} \sum_{c=1}^D (F \cdot F^{-1})_{d_{q_k}, c} \cdot (F \cdot F^{-1})_{d_{p_k}, c} \\ &= \prod_{k=1}^{W^2/2} \sum_{c=1}^D 1_{[d_{q_k}=c]} \cdot 1_{[d_{p_k}=c]} \\ &= \prod_{k=1}^{W^2/2} 1_{[d_{q_k}=d_{p_k}]} \end{aligned}$$

which means $\llbracket \mathcal{A}(\mathbf{h}) \rrbracket_{(I,J)}$ equals to the Kronecker product of $\frac{W^2}{2}$ $M \times M$ -identity matrices, up to permutations of its rows and columns which do not affect its matrix rank. Thus, $\text{rank}(\llbracket \mathcal{A}(\mathbf{h}) \rrbracket) = M^{\frac{W^2}{2}}$.

For the “shared” setting, we denote the parameters of the first GC layer by $\{(\mathbf{w}^{(d)}, \mathbf{b}^{(d)})\}_{d=1}^D$, and set them as:

$$w_{m,j,i}^{(d)} = \begin{cases} (F^{-1})_{m,c} & (\exists c \in [M] \exists u, v \in [W], d = cW^2 + uW + v) \\ & \text{and } \left(\exists k \in \left[\frac{W^2}{2} \right], q_k = (u, v) \text{ and } (j, i) \in \{(1, 1), p_k - q_k + (1, 1)\} \right) \\ 0 & \text{Otherwise} \end{cases}$$

$$b_{j,i}^{(d)} = \begin{cases} 0 & (\exists c \in [M] \exists u, v \in [W], d = cW^2 + uW + v) \\ & \text{and } \left(\exists k \in \left[\frac{W^2}{2} \right], q_k = (u, v) \text{ and } (j, i) \in \{(1, 1), p_k - q_k + (1, 1)\} \right) \\ 1_{[d \leq MW^2]} & \text{Otherwise} \end{cases}$$

the parameters of the other layers are set as in the “unshared” case, and the proof follows similarly. \square

In the preceding claim, we have found a separate example for each matricization, such that the matricization rank is exponential. In the following proof of the theorem 2, we leverage basic properties from measure theory to show that almost everywhere the induced grid tensor has an exponential matricization rank, under every possible even matricization – without explicitly constructing such an example.

Proof. (of theorem 2) For any even partition (I, J) of $\{(1, 1), \dots, (W, W)\}$, according to claim 8 there exist parameters for which $\text{rank}(\llbracket \mathcal{A}(\mathbf{h}) \rrbracket)_{(I, J)} = M^{\frac{W^2}{2}}$, and thus according to lemma 3 the set of parameters for which $\text{rank}(\llbracket \mathcal{A}(\mathbf{h}) \rrbracket) < M^{\frac{W^2}{2}}$ is of measure zero. Since the finite union of sets of measure zero is also of measure zero, then almost everywhere the parameters results in networks such that for all even partitions $\text{rank}(\llbracket \mathcal{A}(\mathbf{h}) \rrbracket)_{(I, J)} = M^{\frac{W^2}{2}}$. Applying lemma 2 concludes the proof. \square